

Abnormal Behavior Detection Finds Malicious Actors

A Cyber Security Assumption Buster Workshop Series

Assertion: “Abnormal behavior detection finds malicious actors”

In an effort to reduce losses due to fraud, financial services companies have been fairly successful in establishing fraud detection analytics, based on abnormal behavior identification, which identify financial transactions that seem out of norm for a particular financial services customer. For example, credit card companies acting on this information will contact cardholders to validate anomalous behavior, or if costs are high, and users unavailable, can freeze accounts until the anomaly is investigated. In this way, they can curtail the loss due to prolonged invalid use of a credit card. Fraud detection algorithms (based on user behavior models) and procedures immediately set off account alarms and/or deny additional transactions after they have detected a fraudulent or suspicious transaction. Depending upon the fraud method (e.g., automated gasoline purchase), they may not always block the first fraudulent transaction on a given card.

Online banking financial institutions employ similar behavioral models to monitor the size and destinations of financial transfers and/or on-line transactions (such as change of address or payee) and will delay transfers until the customer can be reached to confirm the transactions and/or provide additional authentication. Despite the use of best available behavior modeling and monitoring, financial institutions continue to sustain significant financial loss from fraud. Can the field of fraud detection (and cybersecurity in general) be improved by new technology and approaches?

Fraud detection works on the assumption that malicious fiscal behavior is a subset of abnormal behavior – if the fraudulent user mimics the financial behavior of the authorized user, these methods do not work. Detection methods do not assume that malicious behavior is automatically distinguishable from unusual behavior on the part of authorized users. The fraud detection algorithms use the financial services customer’s history to build a profile of “normal” transactions and develop thresholds for unusual behavior. The volume of transactions allows for reasonable thresholds to be established. Fraud detection methods rely on strong models of normal behavior, or known criminal behavior characteristics. The development of many of these models is aided by the fact that the value of a transaction is numeric and allows sets of values to be analyzed with well understood algorithms. For example, credit card purchases have relatively small and fixed semantics: store names are typed, businesses are categorized, relationships among businesses and purchases by card users are fairly easy to establish (e.g., people who buy plane tickets may also purchase luggage, or may eat out more when they are away, or may spend more in general while traveling). These models enable gradual change in behavior to be learned and help drive down false alerts.

Many cyber intrusion detection techniques, or insider threat detection techniques, aim to achieve similar results by using abnormal behavior detection as a starting point. Yet, it is an open question whether these techniques can expect to attain the same broad-based success when applied in the broader cyber security domain. The domains share an adversarial dynamic that might indicate that similar analyses could be effective. But do the assumptions of the relationship between malicious and normal behavior hold true? Can we establish a solid footing in terms of models of normal transaction semantics and transaction value? Does the real time nature of cyber decision making, and the ease of dynamic changes in the criminal’s attack signature, present insurmountable challenges for behavioral techniques?

In this workshop, representatives from government and industry financial organizations will present different financial services fraud detection mechanisms, strengths, and areas needing further development. This will allow workshop participants to have a common understanding of the state of fraud detection practice. Submitted position papers should state your opinion of the assertion and explore new ideas to improve fraud detection specifically and malicious cyber behavior in general.

Issues to be addressed:

- Is malicious cyber behavior a subset of abnormal cyber behavior, or is some malicious behavior normal?
- Is (all or some) malicious behavior a subset of abnormal behavior, or can it be distinguished from benign abnormal behavior?

- If malicious behavior cannot be distinguished from benign behavior, what is the down side of false positives? Of false negatives? Can a threshold be established? In general, or in a context-specific way?
- Can we establish solid semantics for cyber interactions? For some subset of them? Can it be efficient enough to be useful in real time decision making?
- Could behavior modeling be improved by integrating data from different sources such as: transaction location, device ID, time of day and week; integrating activity across call center, web, mobile, POS; correlation across multiple entities, and comparison with past known patterns of criminal behavior?
- It is relatively straightforward to establish a financial transaction baseline for individual users and for populations? Can we establish reasonable cyber baselines? Over the whole space? Over certain types of transactions? Over certain types of users?
- Are there privacy issues associated with cyber behavior analysis? Users expect financial institutions to record and protect their customer's financial information but not maintain a model of their private behavior. When does all the aggregated data and analytic capability of a behavior model impact an individual's privacy? What entity could fill the same role in cyber space? How tolerant should people be about the use of their personnel information?
- Criminals spend time planning and testing prior to execution. Can behavioral models be built to provide early warnings and alerts to fraud in the planning stages, and how could this information be best employed to prevent, detect or mitigate the eventual fraud?